PERSPECTIVE

Systemic risk in banking ecosystems

Andrew G. Haldane[1] & Robert M. May[2]

In the run-up to the recent financial crisis, an increasingly elaborate set of financial instruments emerged, intended to optimize returns to individual institutions with seemingly minimal risk. Essentially no attention was given to their possible effects on the stability of the system as a whole. Drawing analogies with the dynamics of ecological food webs and with networks within which infectious diseases spread, we explore the interplay between complexity and stability in deliberately simplified models of financial networks. We suggest some policy lessons that can be drawn from such models, with the explicit aim of minimizing systemic risk.

In the 1960s, the notion of the 'balance of nature' played a significant part as ecologists sought a conceptual foundation for their subject. In particular, Evelyn Hutchinson[1], following Elton[2], suggested that "oscillations observed in arctic and boreal fauna may be due in part to the communities not being sufficiently complex to damp out oscillations". He went on to state, based on a misunderstanding of MacArthur's[3] paper, that there was now a "formal proof of the increase in stability of a community as the number of links in its food web increases".

To the direct contrary, however, a closer examination of model ecosystems showed that a random assembly of N species, each of which had feedback mechanisms that would ensure the population's stability were it alone, showed a sharp transition from overall stability to instability as the number and strength of interactions among species increased. More explicitly, for $N \gg 1$ this transition occurs once $m\alpha^2 > 1$, where m is the average number of links per species, and ($\pm$) α their average strength[4].

In ecology this has, since the 1970s, prompted a search for special food-web structures that may help reconcile complexity with persistence or stability[5–8]. Along these lines there is, for example, tentative evidence for modularity[9] (particularly in plant–pollinator associations, where linkages tend to be overdispersed or disassociative), and more generally for nested hierarchies in food webs[10]. The fact that some features of the network structure of interactions (such as predator/prey ratios) inferred from the Burgess Shale communities are similar to those in present day ones[11] reinforces hopes that this is a meaningful area of research.

In the wake of the global financial crisis that began in 2007, there is increasing recognition of the need to address risk at the systemic level, as distinct from focusing on individual banks[12,13]. This quest to understand the network dynamics of what might be called 'financial ecosystems' has interesting parallels with ecology in the 1970s. Implicit in much economic thinking in general, and financial mathematics in particular, is the notion of a 'general equilibrium'. Elements of this belief underpin, for example, the pricing of complex derivatives. But, as shown below, deeper analysis of such systems reveals explicit analogies with the concept that too much complexity implies instability, which was found earlier in model ecosystems.

There are, of course, major differences between ecosystems and financial systems. For one thing, today's ecosystems are the winnowed survivors of long-lasting evolutionary processes, whereas the evolution of financial systems is a relatively recent phenomenon[14]. Nor have selective pressures been entirely dispassionate, with the hand of government a constant presence shaping financial structures, especially among institutions deemed "too big to fail"[15]. In financial ecosystems, evolutionary forces have often been survival of the fattest rather than the fittest.

In what follows, we first consider the role of the growth in intrafinancial system claims in generating bank failure and instability, focusing on the problems inherent in prevailing methods of pricing complex derivatives, or arbitrage pricing theory (APT). Second, we sketch various ways in which such an initial bank failure, or 'shock', may propagate to cause cascades of subsequent failure. Third, we outline some tentative policy lessons that might be drawn from these deliberately oversimplified models. Last, we ask how we might reshape the financial system to realize the economic benefits individual banks can deliver, while at the same time paying deliberate and explicit attention to their system-wide stability.

Potential causes of an initial shock

Events external to the banking system, such as recessions, major wars, civil unrest or environmental catastrophes, clearly have the potential to depress the value of a bank's assets so severely that the system fails. Although probably exacerbated by such events, including global imbalances (China as producer and saver, the United States as consumer and debtor), the present crisis seems more akin to self-harm caused by overexuberance within the financial sector itself. Perhaps as much as two-thirds of the spectacular growth in banks' balance sheet over recent decades reflected increasing claims within the financial system, rather than with non-financial agents. One key driver of this explosive intrasystem activity came from the growth in derivative markets.

In 2002, when Warren Buffet first expressed his view that "derivatives are financial weapons of mass destruction"[16], markets—although booming—seemed remarkably stable. Their subsequent growth, illustrated in Fig. 1, has been extraordinary, outpacing the growth in world gross domestic product (GDP) by a factor of three. In some derivatives markets, such as credit default swaps (CDS), growth has outpaced Moore's Law. These developments contributed significantly towards an unprecedented influx of mathematically skilled people (quantitative analysts) into the financial/banking industry. These people produced very sophisticated techniques (including APT), which seemingly allowed you to put a price on future risks, and thus to trade increasingly complex derivative contracts—bundles of assets—with risks apparently decreasing as the bundles grew.

However, recent empirical and theoretical studies have indicated that the trading activity associated with derivatives can have significant effects on markets[17–19]. More specifically, Brock and colleagues[20] have shown that proliferation of hedging instruments can destabilize markets. Building on this, Caccioli and colleagues[21] note that APT makes several conventional assumptions upon which everything else depends: "perfect competition, market liquidity, no-arbitrage and market completeness". Crucially, this adds up to the implicit assumption that trading activity has no feedback on the dynamical behaviour of markets. And indeed, in the APT-fuelled

[1]Bank of England, Threadneedle Street, London EC2R 8AH, UK. [2]Zoology Department, Oxford University, Oxford OX1 3PS, UK.
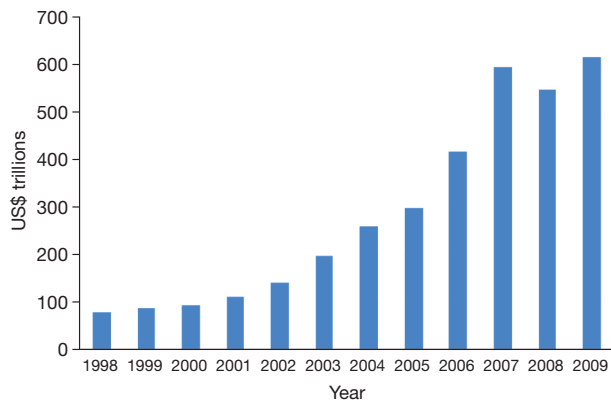
Figure 1 | Notional principal value of outstanding derivative contracts, as recorded at year end. These include foreign exchange, interest rates, equities, commodities and credit derivatives. Data from UK Department for Business, Innovation and Skills, International Monetary Fund and Bank of England calculations.

boom time that preceded the bust, APT seemed to be very successful. In its imaginary world, market failures are caused by regulatory carelessness, resulting in a focus on creating institutional arrangements that seek to guarantee the premises upon which APT is based[22]. To the contrary, Caccioli and colleagues argued[21] that APT is not a 'theory' in the sense habitually used in the sciences, but rather a set of idealized assumptions on which financial engineering is based; that is, APT is part of the problem itself.

Caccioli and colleagues[21] illustrate their point by exploring the dynamical properties of a model that gives a more realistic caricature of markets, going beyond the idealized world of APT to include the effects of individual trades on prices. Prices now depend on the balance between demand and supply. The outcome is that "the road to efficient, arbitrage-free, complete markets can be plagued by singularities which arise upon increasing financial complexity"[21].

Figure 2 illustrates the main results of the analysis by Caccioli and colleagues[21]. Here n is essentially a measure of the proliferation of derivatives or similar financial instruments, and s is the overall average value of the supply of any one such derivative/financial instrument. The parameter ε encodes the risk premium that banks require for trading derivatives[21].

We see from Fig. 2 that if n is less than n^* (here $n^* = 4.14$), the average supply of derivatives, s, is relatively steady and essentially independent of the banks' risk premium (as measured by ε). But as market complexity increases, so that n approaches n^*, there is a sharp singularity at $\varepsilon = 0$.
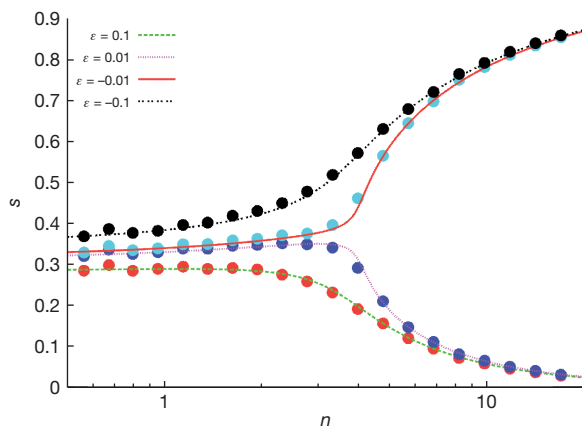

Figure 2 | Discontinuous transition to instability of derivatives as complexity increases. Average supply of any one derivative, s, at competitive equilibrium as a function of the number, n, of different derivatives being traded, for various values of banks' risk premium, ε. Adapted with permission from ref. 21. For fuller discussion, see text.

For $n > n^*$, the average supply increases with increasing n (that is, increasing proliferation of financial instruments) if $\varepsilon > 0$. Conversely, for $\varepsilon < 0$ the supply decreases with increasing complexity once $n > n^*$. It is emphasized[21] that such sensitivity in market behaviour in the neighbourhood of the singularity can easily produce very strong fluctuations—either positive or negative—in the volume of trading in derivative markets.

Note that the consequences of this singularity are not easily intuited from the competitive equilibrium setting. It seems to us that the basic process—in grossly simplified terms—is that once there are enough derivatives to span the space of available states of nature (the net supply of derivatives within the system necessary to meet true hedging demand from non-banks), the market is essentially complete in the sense of the Arrow–Debreu[23] model. Once that happens, gross derivatives positions within the system are essentially unbounded. So long as there is an incentive to supply new instruments—a positive premium to trading—banks will continue to expand gross positions, independent of true hedging demand from non-banks. Such trades are essentially redundant, increasing the dimensionality and complexity of the network at a cost in terms of stability, with no welfare gain because market completeness has already been achieved.

Caccioli and colleagues[21] also examine a measure of market volatility as the risk premium parameter ε varies. If they calculate this quantity under the approximation that the fluctuations in the values of the individual 'supply variables' (s_i; derivatives, etc) are completely uncorrelated, they in effect recover the happy world of APT, with no singularities. This strongly indicates that the highly important singularities in their accurate and self-consistent calculations, with market dynamics included, are associated with the supplies of different derivatives being strongly correlated in this domain, as has found to be the case among derivatives markets in practice.

In summary, Caccioli and colleagues suggest that the idealized assumptions upon which recent financial engineering has been based can give a misleading account of potential instabilities in markets. They also note that these instabilities echo those that can develop in ecosystems as complexity increases[4,24].

Propagation of shocks within financial systems

In ecology's models of food webs, aimed at qualitative understanding of their dynamical response to perturbation, the nodes are simply species, linked to other nodes/species as prey, predator, competitor or mutualist. In epidemiological networks, the nodes are susceptible, infected/infectious or recovered/immune individuals linked by sexual or other contacts. But in a minimally realistic caricature of financial networks—henceforth called banks—the nodes have a more complex structure.

Following Nier and colleagues[25] and Gai and Kapadia[26], we define such a bank/node as schematically illustrated in Fig. 3. In this deliberately oversimplified scheme, a bank's activities are partitioned among four categories. Two represent assets: interbank loans (l_i) and external assets (e_i). The other two represent liabilities: interbank borrowing (b_i) and deposits (d_i). The subscript i labels the specific bank (i = 1, 2, ..., N for a total of N banks). Solvency requires that the difference between a bank's assets and its liabilities (the capital reserve or 'net worth', labelled γ_i in Fig. 3) be positive. That is, $\gamma_I \equiv (e_i + l_i) - (d_i + b_i) \geq 0$.

These banks are now assumed to be interlinked in a random, Erdős–Rényi network, with any one of the N banks connected to any other as lender or borrower, or possibly both, each with probability p. A bank's average number of incoming/borrowing or outgoing/lending links is then $z = p(N - 1)$.

Various further assumptions are now made to carry these Bank of England/Federal Reserve Bank of New York models to the point where the knock-on effects of a single bank failure can be explored in numerical simulations. Much of the essential findings of such studies can be captured, and made more transparent, by a 'mean-field' approximation in which each bank has exactly average behaviour[27]. This means all banks are the same size (rescaled to 1), every bank is linked to exactly z others, all loans
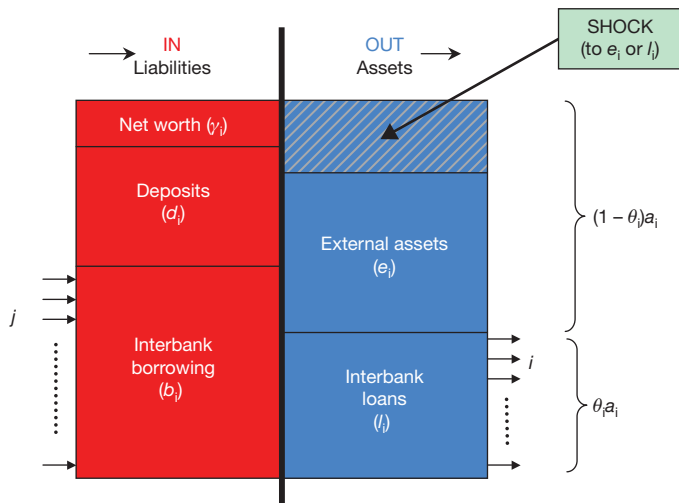
Figure 3 | Schematic model for a node in the interbank network. Adapted with permission from ref. 25.

have the same magnitude, w, as do the capital reserves, γ, and the ratios of loans to total assets, θ.

As illustrated in Fig. 3, all these models study the consequences of a shock that initially hits a single bank, wiping out a fraction, f, of its external assets. If the magnitude of this shock exceeds the capital reserve, $f(1 - \theta) > \gamma$, the bank fails. This is a deliberate oversimplification, aimed at a clearer understanding of how an initial failure can propagate shocks throughout the system.

The most direct effect of such a failure is that its z creditor banks will lose part or all of their loans. If such losses exceed γ, these banks in turn will fail, propagating a third phase of shocks to those remaining, and so on. Note, however, that a failing bank's losses are in effect divided among its z creditors, so that each subsequent phase of loan-driven shocks is attenuated, approximately by a factor z.

Figure 4 illustrates one of the tentative messages emerging from this toy model, showing regimes of failure in terms of the critical parameter γ (capital reserves relative to bank size) and θ (interbank activity as a fraction of total assets). Within the unhatched triangle (0, 1, f), the initially shocked bank fails; in the blue triangle (0, 1, A) a second tranche
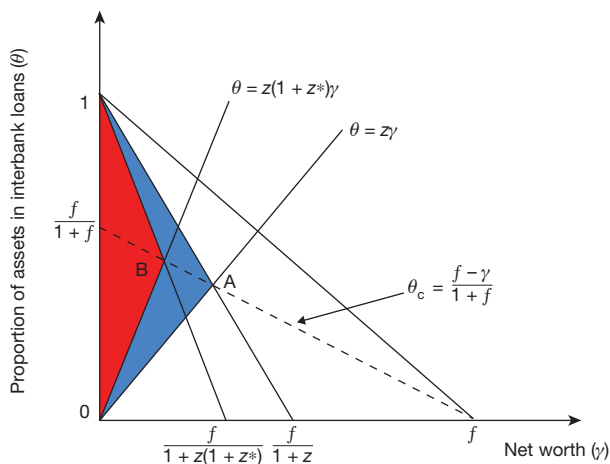

Figure 4 | Domains of interbank lending. Domains are expressed as a fraction of total assets, θ, and capital reserves or net worth, γ, which result in the propagation of interbank loan shocks. The triangle (1, 0, f) defines the region where loss of a fraction f of a bank's external assets will cause it to fail. The blue triangle (1, 0, A) depicts the region in which creditors of the initially failing bank will receive phase II shocks which cause them also to fail, and the red area (1, 0, B) shows the region in which phase III shocks cause failure. Adapted with permission from ref. 27.

of z banks go down; in the red triangle (0, 1, B) there is a third phase of roughly z^2 failures; and so on. Note that, when (as above) the initial shock is to external assets, the system's fragility is maximized (failures for relatively large values of γ) by θ having values intermediate between 0 and 1, which in some ways very roughly corresponds to banks substantially engaged in both retail and investment (high-street and casino) activity. As seen earlier and in Fig. 4, an increase in the system's connectivity, z, causes the coloured region of instability to shrink; high connectivity distributes, and thereby attenuates risk. On the other hand, when later-phase failures do occur, they will then involve more banks.

A second, and almost surely more important, source of shock propagation arises from losses in the value of a bank's external assets, caused by a generalized fall in market prices, a rise in expected defaults or a failing bank's 'fire sale' actions. Such market liquidity shocks are conventionally and sensibly represented by discount factors that, for a given asset class, are proportional to the number of failing banks holding the asset. This may be generalized to distinguish between strong liquidity shocks, associated with discounting specific asset classes, and weak liquidity shocks, resulting from the expectation of further defaults or a more general loss of confidence[27]. In all cases and in sharp contrast to the attenuation in interbank loan shocks, liquidity shocks amplify as more banks fail. Thus, relatively small initial liquidity shocks have the potential to make strong contributions to systemic risk.

A third mechanism of shock propagation, which has been a marked—and in many peoples' opinion the most important—feature of the recent crisis has been the diminished availability of interbank loans, or in the jargon of the trade, 'funding liquidity shocks'. This has often taken the form of liquidity hoarding in interbank funding markets. Gai and Kapadia[28] have recently shown how such liquidity hoarding can cascade through a banking network, with severe consequences. As one bank calls in or shortens the term of its interbank loans, affected banks tend in turn to do the same. The result is a liquidity-hoarding shock that is not subject to the attenuation characteristic of interbank default shocks.

All three propagation mechanisms can be drawn together within the framework defined by Fig. 3 (see also N. Arinaminpathy, S. Kapadia and R.M.M., manuscript in preparation). The model can also be generalized to treat banks of varying size, including the extreme but realistic case of a few very large all-purpose banks, each connected to many smaller banks; interconnectivity within real banking networks is far from random[29–31], with long-tailed degree distributions. It also seems that these networks tend to be disassociative rather than proportionally connected: that is, big banks are disproportionately linked to smaller ones, and conversely. Such a 'wiring up' of a network is known, unfortunately, to maximize the number of individuals infected by an agent that is transmitted by interpersonal contact[32]. On the other hand, such disassociative structures are likely to support a larger number of coexisting banks (another link between ecology and banking[33]), and can make the network more robust to random losses[9,34].

Some of this work, particularly that on liquidity shocks, echoes an important insight from pervious work[35,36] (N. Beale and colleagues, manuscript in preparation). This is that excessive homogeneity within a financial system—all the banks doing the same thing—can minimize risk for each individual bank, but maximize the probability of the entire system collapsing. A very simple toy model illustrates this. Suppose you have N banks and N distinct, uncorrelated asset classes, each of which has some very small probability, ε, of having its value decline to the extent that a bank holding solely that asset would fail. At the inhomogeneous extreme, assume each bank holds the entirety of one of the N assets: the probability for any one bank to fail is now ε, whereas that for the system is a vastly smaller ε^N. At the opposite, homogeneous extreme, assume all banks are identical, each holding $1/N$ of every one of the N assets: the probability for any one bank to fail can now be calculated as $N^N \varepsilon^N / N!$, and this is obviously also the probability for all N of these banks to fail. This homogeneous, 'herding behaviour' limit clearly makes each individual bank safer, but the systemic risk is much larger. More realistic versions of this scenario consistently show the same unhappy

conclusion. Tentative evidence comes from the fact that the world's five largest banks have shown increasing concentrations of assets over the last ten years, in contrast to the top five hedge funds, whose less concentrated systems can give greater scope for diversity. The former are in trouble, the latter much less so.

Implications for public policy

All the studies described earlier involve numerical simulations, but many combine such work with analytic results of the kind exemplified by Fig. 4. Such analysis of the dynamics of deliberately oversimplified models of financial ecosystems carries potentially far-reaching implications for the design and implementation of public policy. These implications include the following.

Setting regulatory capital/liquidity ratios

The cornerstone of the current international regulatory agenda is the setting of higher requirements for banks' capital and liquid assets. The traditional rationale for such requirements is that they reduce idiosyncratic risks to the balance sheets of individual banks. An alternative and more far-reaching interpretation is that they are a means of strengthening the financial system as a whole by limiting the potential for network spillovers. With this wider objective, prudential regulation is following in the footsteps of ecology, which has increasingly drawn on a system-wide perspective when promoting and managing ecosystem resilience.

The systemic rationale for financial regulatory intervention is well illustrated by the dynamic models outlined earlier. Consider banks' buffers of capital or net worth (γ). These capital ratios have been in secular decline in relation to banks' total assets for at least the past 150 years in the United Kingdom and United States[37]. Reversing these trends by setting higher required capital ratios strengthens the absorptive capacity of each of the nodes in the financial network in response to external shocks. As importantly, however, it also lessens the risk of idiosyncratic defaults cascading around the system, as illustrated in Fig. 4.

Broadly, the same arguments apply in the setting of regulatory requirements on banks' liquid assets. These liquidity ratios have also been in secular decline in the United Kingdom and United States, for at least the past half century. Typically, liquidity requirements are specified as a minimum ratio of banks' liquid assets to their short-term liabilities. This liquidity ratio can be seen as a means of short-circuiting the potential for systemic liquidity spillovers arising from fire sales on the asset side of the balance sheet (liquidity shocks) or liquidity hoarding on the liabilities side (liquidity-hoarding shocks). In particular, holdings of liquid assets reduce the potential for market liquidity risk to propagate around the system, while limits on short-term liabilities reduce the spread of funding liquidity risk around the system.

Setting systemic regulatory requirements

Looking at financial risk through a network lens indicates a fundamentally different rationale for prudential regulation. It also indicates a quite different calibration of such regulation. Prudential regulation has become increasingly risk-based with the advent of first Basel I and latterly Basel II. But the risk in question to which regulation was then calibrated has tended to be institution-specific rather than systemic risk.

To give an example, as conventionally calibrated, capital regulation seeks to equalize failure probabilities across individual institutions to a given tolerance threshold—such as a 0.1% probability of failure. Approaching this problem from a system-wide angle indicates a rather different calibration. Instead, the objective would be to set firms' capital requirements to equalize the marginal cost to the system as a whole of their failure. In other words, regulatory requirements would be set higher for those banks bringing greatest risk to the system; for example, because of their size or connectivity.

Although new in the context of banking, the essential insight here is an old one in the study of epidemiological networks. Anderson and May[38] established the theoretical case for focusing preventative action

on 'super-spreaders' within the network to limit the potential for system-wide spread. Although initially applied in the study of contagious diseases, such as HIV/AIDs, this same insight has since been applied in managing the dynamics of the world wide web, power grids and biological ecosystems[8,39].

If anything, this same logic applies with even greater force in banking. There has been a spectacular rise in the size and concentration of the financial system over the past two decades, with the rapid emergence of 'super-spreader institutions' too big, connected or important to fail (Fig. 5). The collateral damage, to both the real economy and financial system, following the failure of Lehman Brothers in October 2008 is testimony to the force of such super-spreader dynamics. Protecting the financial system from future such events would require the key super-spreader nodes to run with higher—potentially much higher—buffers of capital and liquid assets, which are then proportional to the system-wide risk they contribute.

A second source of system-wide risk, in addition to super-spreader failures, arises from aggregate external events, such as booms and busts in the real economy. Indeed, historically this has been the largest single source of banking problems. If regulation could be operated counter-cyclically, with buffers rising in booms and falling in recessions, this would lessen systemic risk from this particular source. Why? Because increasing insurance in a boom would increase system-wide resilience against the subsequent bust, as well as providing an incentive for banks to curb risk-taking during the boom. Operating regulation in this way would be a new departure for prudential policy—so-called macro-prudential policy—but a potentially important one[40,41] from a systemic risk perspective.

Netting and clearing derivatives

The rapid growth in the size and complexity of the derivatives market contributed importantly to the destabilizing dynamics of the system under stress during the recent financial crisis. This begs questions about the underlying structure and dimensionality of the derivatives market. One means of simplifying the complex web of interactions between banks in derivatives markets is to centralize the trading and clearing of these instruments. For example, central counterparties interpose themselves between every bilateral transaction, thereby replacing a cat's-cradle of financial network interactions with a single hub-and-spokes configuration. Provided the central counterparty is extremely robust—to prevent it becoming a super-spreader itself—the upshot is
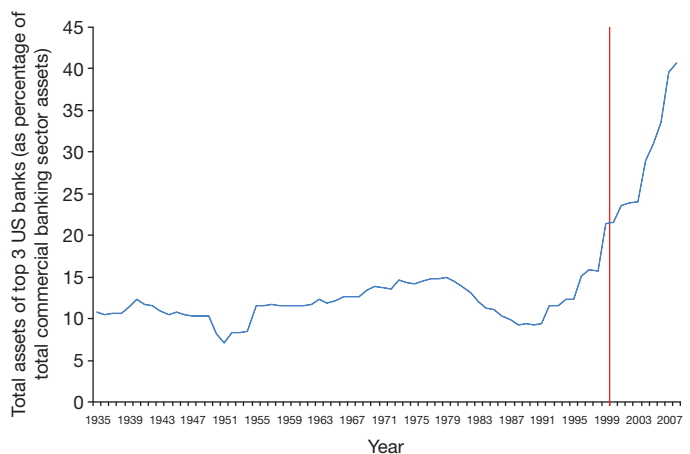

Figure 5 | Recent rise in the size and concentration of the United States financial system. This figure illustrates the marked increase in asset concentration within the United States banking system since the Glass–Steagal restrictions were revoked in 1999. Red line represents the Gramm–Leach–Bliley Act (1999), which revoked Glass–Steagall restrictions. Data include only the insured depository subsidiaries of banks to ensure consistency over time; for example, non-deposit subsidiaries are not included. Data from the Federal Deposit Insurance Corporation.

a less complex and lower-risk financial network. Efforts are underway internationally to extend the scope and reach of central counterparty clearing, in particular to ensure it covers transactions in complex over-the-counter derivative instruments, such as CDS (D. Duffie and H. Zhu, manuscript in preparation).

In parallel, there are international efforts to reduce the dimensionality of derivatives contracts by eliminating redundant trades and through netting. This redundancy might arise either because contracts have been reassigned to participants (but the claim not extinguished) or because there are perfectly offsetting bilateral transactions between two parties that can be netted. For example, the stock of CDS contracts has already been reduced by around $25 trillion since December 2007 as a result of such netting arrangements. Looking forward, there may be more sophisticated multilateral netting algorithms that can be used to reduce further derivatives balances.

Shaping the topology of the financial network

The analytic model outlined earlier demonstrates that the topology of the financial sector's balance sheet has fundamental implications for the state and dynamics of systemic risk. From a public policy perspective, two topological features are key[15].

First, diversity across the financial system. In the run-up to the crisis, and in the pursuit of diversification, banks' balance sheets and risk management systems became increasingly homogenous. For example, banks became increasingly reliant on wholesale funding on the liabilities side of the balance sheet; in structured credit on the assets side of their balance sheet; and managed the resulting risks using the same value-at-risk models. This desire for diversification was individually rational from a risk perspective. But it came at the expense of lower diversity across the system as whole, thereby increasing systemic risk. Homogeneity bred fragility (N. Beale and colleagues, manuscript in preparation).

In regulating the financial system, little effort has as yet been put into assessing the system-wide characteristics of the network, such as the diversity of its aggregate balance sheet and risk management models. Even less effort has been put into providing regulatory incentives to promote diversity of balance sheet structures, business models and risk management systems. In rebuilding and maintaining the financial system, this systemic diversity objective should probably be given much greater prominence by the regulatory community.

Second, modularity within the financial system. The structure of many non-financial networks is explicitly and intentionally modular. This includes the design of personal computers and the world wide web and the management of forests and utility grids. Modular configurations prevent contagion infecting the whole network in the event of nodal failure. By limiting the potential for cascades, modularity protects the systemic resilience of both natural and constructed networks.

The same principles apply in banking. That is why there is an ongoing debate on the merits of splitting banks, either to limit their size (to curtail the strength of cascades following failure) or to limit their activities (to curtail the potential for cross-contamination within firms). The recently proposed Volcker rule in the United States, quarantining risky hedge fund, private equity and proprietary trading activity from other areas of banking business, is one example of modularity in practice. In the United Kingdom, the new government have recently set up a Royal Commission to investigate the case for encouraging modularity and diversity in banking ecosystems, as a means of buttressing systemic resilience.

It took a generation for ecological models to adapt. The same is likely to be true of banking and finance.

1. Hutchinson, G. E. Homage to Santa Rosalia, or why are there so many kinds of animals? *Am. Nat.* **93,** 145–159 (1959).
2. Elton, C. S. *The Ecology of Invasions by Animals and Plants* (Methuen, 1958).
3. MacArthur, R. H. Fluctuations of animal populations, and a measure of community stability. *Ecology* **36,** 533–536 (1955).
4. May, R. M. Will a large complex system be stable? *Nature* **238,** 413–414 (1972).
5. Stouffer, D. B. *et al.* Quantitative patterns in the structure of model and empirical food webs. *Ecology* **86,** 1301–1311 (2005).
6. Pascuale, M. & Dunne, J. A. *Ecological Networks: Linking Structure to Dynamics in Food Webs* (Oxford Univ. Press, 2006).
7. Dunne, J. A. *et al.* Network structure and robustness of marine food webs. *Mar. Ecol. Prog. Ser.* **273,** 291–302 (2004).
8. May, R. M. Network structure and the biology of populations. *Trends Ecol. Evol.* **21,** 394–399 (2006).
9. Bascompte, J. Disentangling the web of life. *Science* **325,** 416–419 (2009).
10. Sugihara, G. & Ye, H. Cooperative network dynamics. *Nature* **458,** 979–980 (2009).
11. Dunne, J. A. *et al.* Compilation and network analyses of Cambrian food webs. *PLoS Biol.* **6,** e102 (2008).
12. Haldane, A. G. Rethinking the financial network. 〈http://www.bankofengland.co.uk/publications/speeches/2009/speech386.pdf〉 (2009).
13. Jones, C. Preventing system failure. *Cent. Banking* **21,** 69–75 (2010).
    The survey of recent thinking about financial networks, in lay language.
14. Farmer, J. D. Market force, ecology and evolution. *Ind. Corp. Change* **11,** 895–953 (2002).
15. Haldane, A. G. The $100 billion question. 〈http://www.bankofengland.co.uk/publications/speeches/2010/speech433.pdf〉 (2010).
16. Buffet, W. E. Chairman's Letter. *Berkshire Hathaway Inc. 2002 Annual Report* 15 (2002).
    A prescient and clearly expressed warning of problems inherent in derivatives.
17. Sircar, K. R. & Papanicolau, G. General Black-Scholes models accounting for increased market volatility from hedging strategies. *Appl. Math. Finance* **5,** 45–82 (1998).
18. Avellaneda, M. & Lipkin, M. D. A market-induced mechanism for stock pinning. *Quantit. Finance* **3,** 417–425 (2003).
19. Osler, C. L. Macro lessons from microstructure. *Int. J. Finance Econ.* **11,** 55–80 (2006).
20. Brock, W. A., Hommes, C. H. & Wagner, F. O. O. More hedging instruments may destabilise markets. *J. Econ. Dynam. Cont.* **33,** 1912–1928 (2008).
21. Caccioli, F., Marsili, M. & Vivo, P. Eroding market stability by proliferation of financial instruments. *Eur. Phys. J. B* **71,** 467–479 (2009).
    A sophisticated and important analysis of a major flaw in the pricing of derivatives.
22. Pliska, S. R. *Introduction to Mathematical Finance: Discrete Time Models* (Blackwell, 1997).
23. Arrow, K. J. & Debreu, G. Existence of an equilibrium for a competitive economy. *Econometrica* **22,** 265–290 (1954).
24. May, R. M. *Stability and Complexity in Model Ecosystems* (Princeton Univ. Press, 1973).
25. Nier, E., Yang, J., Yorulmazer, T. & Alentorn, A. Network models and financial stability. *J. Econ. Dyn. Control* **31,** 2033–2060 (2007).
26. Gai, P. & Kapadia, S. Contagion in financial networks. *Proc. R. Soc. A* **466,** 2401–2423 (2010).
27. May, R. M. & Arinaminpathy, N. Systemic risk: the dynamics of model banking systems. *J. R. Soc. Interface* **7,** 823–838 (2010).
28. Gai, P. & Kapadia, S. Liquidity hoarding, network externalities, and interbank market collapse. *Proc. R. Soc. A* **466,** 2401–2423 (2010).
29. Schweitzer, F. *et al.* Economic networks: the new challenges. *Science* **325,** 422–425 (2009).
30. May, R. M., Levin, S. A. & Sugihara, G. Complex systems: ecology for bankers. *Nature* **451,** 893–895 (2008).
    A brief overview of the 2006 NAS/FRBNY study on systemic risk.
31. Kyriakopoulos, F. *et al.* Network and eigenvalue analysis of financial transaction networks. *Eur. Phys. J. B* **71,** 523–531 (2009).
32. Gupta, S., Anderson, R. M. & May, R. M. Networks of sexual contacts: implications for the pattern of spread of HIV. *AIDS* **3,** 807–818 (1989).
33. Bastolla, U. *et al.* The architecture of mutualistic networks minimizes competition and increases biodiversity. *Nature* **458,** 1018–1020 (2009).
34. Memmott, J. *et al.* Tolerance of pollination networks to species extinctions. *Proc. R. Soc. Lond. B* **271,** 2605–2611 (2004).
35. Battiston, A., Gatti, D. D., Gallegati, M., Greenwald, B. C. & Stiglitz, J. E. Liaisons dangereuses: increasing connectivity, risk sharing and systemic risk. 〈http://www.nber.org/papers/w15611〉 (2009).
36. Stiglitz, J. Contagion, liberalization, and the optimal structure of globalization. *J. Global. Develop.* (in the press).
37. Haldane, A. G. Banking on the state. 〈http://www.bankofengland.co.uk/publications/speeches/2009/speech409.pdf〉 (2009).
    This gives a more full account of the policy issues discussed here.
38. Anderson, R. M. & May, R. M. *Infectious Diseases of Humans: Transmission and Control* Ch.12.3 (Oxford Univ. Press, 1991).
39. Barabesi, A.-L. *Linked: The New Science of Networks* (Perseus, 2002).
40. Bank of England. The role of macroprudential policy: a discussion paper. 〈http://www.bankofengland.co.uk/publications/other/financialstability/roleofmacroprudentialpolicy091121.pdf〉 (2009).
41. Turner, P. The debate on financial system resilience: macroprudential instruments. 〈http://www.bankofengland.co.uk/publications/speeches/2009/speech407.pdf〉 (2009).